arXiv:1408.2288v1 [cs.NE] 11 Aug 2014

Genetic Programming for Smart Phone Personalisation

Philip Valencia, Aiden Haak, Alban Cotillon, and Raja Jurdak

Autonomous Systems Laboratory
CSIRO Computational Informatics, Brisbane Australia

Abstract

Personalisation in smart phones requires adaptability to dynamic context based on user mobility, application usage and sensor inputs. Current personalisation approaches, which rely on static logic that is developed a priori, do not provide sufficient adaptability to dynamic and unexpected context. This paper proposes genetic programming (GP), which can evolve program logic in realtime, as an online learning method to deal with the highly dynamic context in smart phone personalisation. We introduce the concept of collaborative smart phone personalisation through the GP Island Model, in order to exploit shared context among co-located phone users and reduce convergence time. We implement these concepts on real smartphones to demonstrate the capability of personalisation through GP and to explore the benefits of the Island Model. Our empirical evaluations on two example applications confirm that the Island Model can reduce convergence time by up to two-thirds over standalone GP personalisation.

Keywords:

1. Introduction

Smartphones have experienced exponential growth in recent years. These phones embed a growing diversity of sensors, such as gyroscopes, accelerometers, Global Positioning Systems (GPS), and cameras, with broad applicability in areas such as urban sensing or environmental monitoring. Coupled with phone users' high mobility and diverse profiles [2], this sensory richness has fuelled complex applications with composite logic, including features such as location-based and usage-based services. The increased application complexity involves significant challenges in personalising smart phone applications so that they can adapt to new or unexpected context.

Smartphone personalisation either occurs centrally at a server or locally at the phone. Centralised approaches track user activity to customise content delivery or application behaviour. They can easily misrepresent user preferences due to spurious activity, and they involve privacy concerns as users need to share their data with the content providers. Most standalone smartphone algorithms, aiming at either data-centric [5] or user-centric personalisation [3], are based on static or rule-based approaches. However, personalisation increasingly depends on contextual information and user inputs [4], which are both subject to dynamic changes arising from mobility and user preferences. The problem of personalisation of smart phones is therefore multidimensional and requires an approach that can not only adapt parameters in response to changes in user preferences and context, but also adapt the program logic to optimise for unpredictable changes in context.

Online learning is well-suited for smart phone personalisation. In order to provide maximum versatility to support the creativity and unpredictability

of smart phone users, we need an online learning approach that allows for the adaptation of program logic, and not just parameters within fixed program logic. Additionally, smart phone users are often co-located with several other users that share their context, providing opportunities for collaborative personalisation based on this shared context. The most suitable online learning strategy should provide the syntactic richness to evolve logic on a single smartphone, and to support the sharing of functional logical blocks among co-located phones according to the building block hypothesis [34].

Most online learning methods, such as reinforcement learning [6] and neural networks [7], are concerned with parameter optimisation only. While these methods could be run within an evolutionary framework, any sharing of genetic sequences across phones would not necessarily map to functional logic blocks, which could slow down convergence. Genetic Programming (GP), on the other hand, is amenable to this scenario, as it evolves both logic and parameters simultaneously, providing it with the syntactic richness and flexibility that is comparable to offline human software development. Because GP evolves functional logic blocks, sharing logic across multiple phones is both more meaningful and more conducive to quicker convergence.

This paper proposes GP for smart phone personalisation. We first show empirically that GP can support smart phone personalisation through our software framework called the Android Genetic Programming Framework (AGP). In order to expedite convergence towards high performing applications, we propose collaborative personalisation of smartphones through the GP Island Model to exploit shared context between co-located phones. The main research question we aim to answer is: "To what extent does collabo-

rative smartphone personalisation improve convergence time?". To address this question, we extend AGP to support the island model and run extensive experiments with 2 Android phones for a comparative evaluation of the benefits of sharing logic among smartphones against *stand-alone* personalisation, exploring 6 scenarios with different migration rates and intensities. Our results show that the Island Model consistently outperforms the standalone GP for online personaliation, and somewhat surprisingly, that injecting random programs into the subpopulations can be beneficial for the more complex application of energy-efficient localisation.

The contributions of this paper are:

- Introduction, motivation, and demonstration of genetic programming for smart phone personalisation
- Proposal of collaborative smart phone personalisation through the GP Island Model for faster convergence and exploitation of shared context
- Empirical evaluation of both standalone and collaborative personalisation through two case study applications, which confirm the benefits of collaborative personalisation

The remainder of the paper is organised as follows. Section 2 discusses related work in the literature. Section 3 defines the problem for smart phone personalisation, motivates genetic programming to address this problem, and proposes collaborative GP for faster personalisation. Section 4 briefly introduces the AGP framework and our extension to support the Island Model. Section 5 demonstrates online personalisation through GP, while Section 6

evaluates the benefits of collaborative personalisation through the Island Model. Section 7 discusses the results and concludes the paper.

2. Related Work

Most online learning approaches, including neural networks, adaptive systems, and reinforcement learning, use specialised structures for representation within the learning process [8]. Because of their reliance on specialised structures, these approaches are amenable for online learning situations where the overall program logic is well-defined while individual parameters within this program structure need to be optimised. While these approaches have been used in an evolutionary context (for e.g. in [31]), their representation of genetic material is not necessarily aligned to functional logic blocks, which makes them less amenable to logic sharing among islands. In contrast, genetic programming uses program logic representation in the learning process, supporting high versatility for entirely new actions in response to unpredictable stimuli. Through its program tree representation, GP can isolate functional logic blocks for sharing across islands to leverage shared context for faster convergence.

Several previous GP works in have focused on architectural issues. Whereas some solutions provide generic frameworks for evolutionary computation problems [10, 11, 9], others propose application-specific solutions. Ismail et al. [12] describe a GP framework for extracting a mathematic formula needed for fingerprint matching, whereas, the authors in [13] focus on a genetic programming framework for content-based image retrieval. In [16], Lacerda et al. introduce a framework for associating ads with web pages based on GP.

Valencia et al. [14] study genetic programming for Wireless Sensor Networks and propose the In Situ Distributed Genetic Programming (IDGP) framework. DGPF [17] brings utilities for Master/Slave, Peer-to-Peer, and P2P/MS hybrid distributed search execution. P-Cage [25] introduces and evaluates a complete framework for the execution of genetic programs in a P2P environment. It shows the relevance of using P2P networks scalability to counteract computation limitations.

Design patterns describe the interaction between groups of classes or objects. They concentrate on specific concerns for implementing source code to support program organization. When they are well integrated into a framework, they ensure the goals of extensibility and reuse. Lenaerts and Manderick [15] discuss the construction of an object-oriented Genetic Programming framework using design patterns to increase flexibility and reusability. McPhee et al. [10] extend the latter to Evolutionary Computation (EC). As the solution search space for a problem becomes wider, it leads to a more abstract set of classes. Based on those works, Ventura et al. introduced JCLEC [9], a Java Framework for evolutionary computation. They present a layered architecture and provide a GUI for EC. This paper similarly uses Java for a genetic programming framework, albeit for a more resource constrained smart phone platform.

Smartphone personalisation research has mainly focused on rule-based approaches. Korpipaa et al. [18] introduced a first framework for user customization using context changes as triggers. Their prototype enables the end-user to set up actions on context events such as GUI interactions, RFID tag informations, accelerometer peaks and other data from embedded sen-

sors. Since this first try on Nokia N73 smartphone, similar applications have been published on both Android and Apple applications stores. Onx [19] and Launch Center Pro [20] let users build rules for automating various tasks. These mechanisms still require explicit user involvement in personalisation, which limits their utility to more technology savvy users.

More recently, Interactive Differential Evolution (IDE) has been applied on smart phones [1] as a method to achieve quick image enhancement of photos taken on mobile phones. As with most Interactive Evolutionary Computation (IEC) implementations, IDE encodes parameters of a fixed-logic solution as a vector and optimises these parameters over time, unlike GP which also optimises the logic.

IEC methods include interactive evolution strategy [27], interactive genetic algorithm [30], interactive genetic programming [29], and human-based genetic algorithm [28]. An interactive genetic algorithm (IGA) is defined as a genetic algorithm that uses human evaluation.

The community-based earthquake detection technique, proposed by Faulkner et al [32], highlighted the need to consider distributed context, using accelerometer data from multiple smartphones to detect earthquakes. The distributed nature of the data allowed the detection and filtering out of false positives and negatives due to spurious sensor data. This highlights the importance of distributed context for the collaborative learning method in this paper.

Another instance of related work is activity recognition using accelerometers data [33]. Interestingly, Weiss found that using recognition models tailored specifically to a user outperformed an impersonal model which used

data gathered from multiple users. The personal model also outperformed a hybrid model using a combination of a model based on data from both a specific user and a model based on data gathered from multiple users. However, it must be noted that these models were all developed offline and did not adapt to their users. The work in this paper postulates and shows that collaborative learning can be useful for adapting to shared context for some applications, while other problems such as motion models may indeed be more amenable to personalised models, where standalone GP could learn the most appropriate logic.

In an earlier paper, we introduced the Android Genetic Programming Framework (AGP) [21] as the first genetic programming solution available on smartphones, demonstrating its ability to provide and update context-specific solutions over time. This paper proposes collaborative smart phone personalisation through the GP Island Model and extends AGP to evaluate this concept [1].

3. Smart Phone Personalisation through GP

This section first defines the smart phone personalisation problem, and then motivates online learning, particularly through GP, to address this problem. The final part of the section argues for using collaborative learning through the Island Model to exploit shared context and shorten convergence time.

[1] AGP source code and documentation are publicly available for download at http://sourceforge.net/projects/agpframework/

3.1. Problem Definition

Mobile phone users have always tried to customise their devices, for instance through personalised ring tones. The emergence of smartphones takes the personalisation possibilities to a new level. First of all, smartphones have access to a broad range of Internet data which can be augmented with sensor-based context information. Secondly, the higher computing performance of smart phones enables developers to create novel applications. The combination of processing power, content accessibility and context awareness opens new opportunities for personalisation. Finally, the intrinsic smart phone user mobility dictates a highly dynamic and unpredictable context.

Personalisation mechanisms have been slow to respond to these opportunities, with applications relying largely on static logic designed by the application programmer. These approaches are unable to cope with unexpected changes in context that is not specified in their static logic, highlighting the need for more dynamic personalisation methods.

The main problem that we are interested in this paper is to support continuous personalisation of smart phone applications in response to changes in environmental context and in user preference. Given that these changes are unpredictable at the time of application development, we propose online learning to address this problem.

3.2. Online Learning for Personalisation

Several machine learning approaches, such as reinforcement learning, neural networks, support vector machines, and genetic programming, can support online learning through integration into an evolutionary framework that iterates through their optimisation process. We compare the suitability of

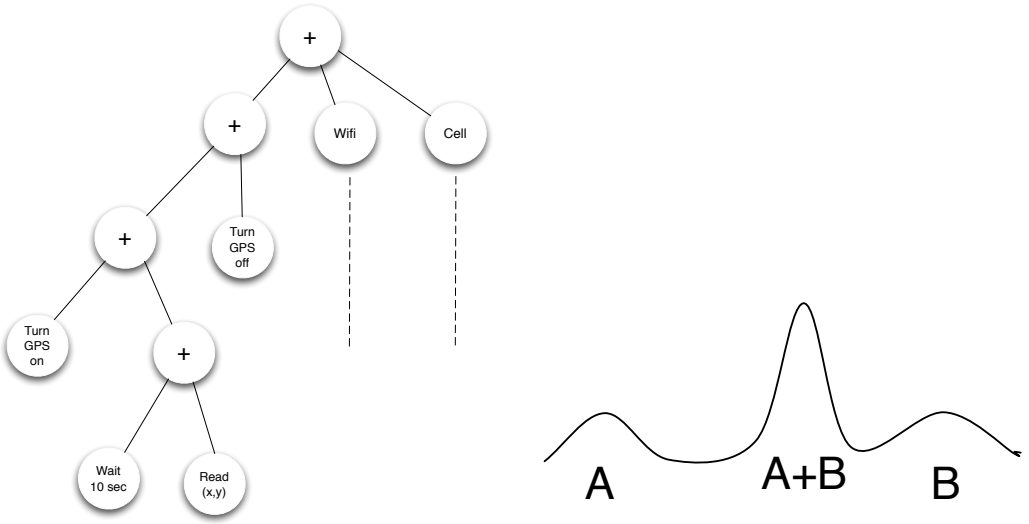

(a) Program tree representation

(b) Combining functional logic blocks to reach peaks in the fitness landscape

Figure 1: Online GP's logic representation supports sharing of functional logic blocks across islands

2 candidate online learning approaches for smart phone personalisation in further detail: (1) neural networks; and (2) genetic programming.

Neural networks typically learn the mapping between a set of input parameters and an output through an offline training process, and the resulting neural network is used for online predictions of new outputs according to new inputs. It is possible to train an existing neural network topology in an online manner; however, such a fixed approach limits the capacity to share blocks of logic. It is also possible to evolve neural networks by representing their topology (and potentially interconnecting weights) as a genetic encoding. This permits the sharing of functional blocks represented by partial networks, yet it is much more likely that functional blocks will be arbitrarily

split, meaning the output of a partial network may not match well as an input to another network. In contrast, online GP ensures through syntactic checking that the output of a block of logic will always be acceptable to another block, which is arguably beneficial in problems where solutions are naturally represented as programs (logical blocks).

Consider the program tree representation in Figure 1(a) which provides a partial view of the energy-efficient localisation application in [21]. The main goal of this application is to dynamically select among various positioning modalities on the smartphone to obtain accurate position estimate subject to energy limitations. The deeper branch in the tree represents learned GPS logic, while the deeper nodes of the other two shallow branches are omitted for simplicity. Let the logic in this GPS branch be functionality A, whose performance in the search space is indicated in Figure 1(b). Assume there is another island that has evolved a separate program with functionality B. Because GP maps functionality to specific branches in the program tree, the sharing of genetic material between the two islands has a reasonable chance of combining the two functionalities A+B to reach much higher performance. A similar merging of such 'blocks of logic' is significantly more challenging with other online learning approaches.

Similarly, when sharing genetic material between subpopulations that have evolved in isolation, the syntactic checking within GP always ensures genetic material can be readily incorporated, unlike evolutionary neural networks where a high level of speciation could occur to the extent where any network from one subpopulation may not be compatible with any network in another subpopulation. Such a scenario would again be detrimental to

collaborative evolution.

3.3. Collaborative GP Personalisation

Personalising complex applications through online learning may require long convergence time and may have only partial contextual information. More complex personalisation may take a long time to converge as the search space grows, consuming a significant portion of the phone's processing cycles and energy. Additionally, sensor data on a smartphones represents a single spatial sample of a physical property, which may not be representative of the broader context to which personalisation is required. For instance, if a mobile user requires personalised application response based on local weather, a smartphone placed in the sun can report ambient temperatures that are much higher than reality, biasing evolution toward non-representative context.

The innate inter-communication capability between mobile devices lends itself to the Island Model implementation where each mobile device hosts a population and evolved programs can be serialised and shared (migrated) [25]. While intuitively one expects the parallel resources will expedite convergence, the Island Model is also known to generate better quality solutions [26].

We argue that distributed logic evolution mitigates the shortcomings of stand-alone online personalisation by sharing genetic material among several smartphones and providing a broader context through increased spatial sampling. Because smartphone users have a tendency to congregate together and possibly form small community-like groups, co-located phones have a significant overlap in their context, which enables quicker and more effective personalisation by sharing logic among these phones.

4. Android Genetic Programming Framework

Faster convergence towards desirable behaviour and learning shared distributed context are both strong motivators for collaborative smart phone personalisation. We choose to demonstrate this approach by extending the AGP framework to support the Island Model. In designing this collaborative learning method, we consider the following:

- Broadcast communication: we select broadcast best-effort communication interfaces for migrating programs to avoid maintaining connection information and to support graceful degradation. When a migrant does not reach its destination, evolution at the destination continues smoothly. Broadcast communication also provides scalability, where a single transmission of a migrant program can be received by multiple co-located phones.

- Flexible configuration: setting parameter values for the collaborative learning, such as the rate at which programs are shared (migration frequency), the number of migrating programs (migration intensity), and the type of migrating programs (mutated or crossover programs), depends on the application. We aim to ensure flexibility in setting these parameters.

- Compatibility: we aim to maintain well-defined application programming interfaces for AGP and to maintain full compatibility with the underlying Android OS.

- Implicit context sharing: collaboration between phones is based on

sharing logic rather than raw data, as the logic on a remote device will have evolved based on the perceived context at that device.

Prior to describing our collaborative smartphone personalisation approach on the basis of the above guidelines, we briefly revisit the Android Genetic Programming (AGP) framework.

4.1. Android Genetic Programming Framework

The AGP framework [21] was developed as an easy-to-use genetic programming framework for evolving multi-objective, context-sensitive, adaptable applications on the Android phone operating system. Figure 2 provides a high level view of the AGP framework. The core AGP framework provides the infrastructure for writing new GP applications for Android smartphones. For full details of the AGP framework the reader is referred to [21]. The main purpose of AGP is to provide online GP infrastructure transparently to users for continuous smart phone personalisation. The original AGP framework supports only standalone GP evolution, which is why we extend it to support the Island Model as Section 4.4 describes.

4.2. Selectors and Genetic Operations

When the evaluation thread is running alongside the execution thread, programs are evaluated and receive a fitness value. In GP, the programs that perform well are chosen to breed the next generation. Selectors are organized following the Strategy Pattern in order to provide flexibility. This way, the developer can easily switch between standard selectors provided by AGP such as the Wheel selector, or create his own selector solution.

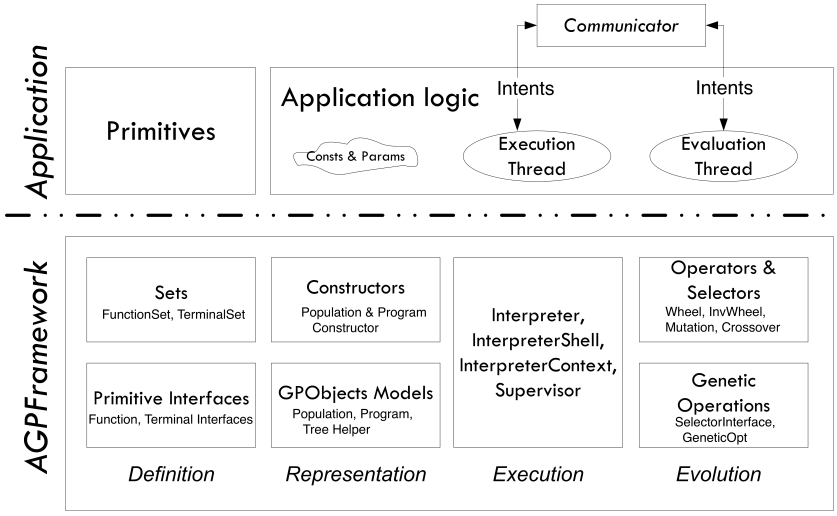

Figure 2: Global Architecture of AGP

AGP then executes genetic operators on selected programs. In this regard, AGP currently supports two primary genetic operators widely used in GP, crossover and mutation, although it is extendible with more operators.

4.3. Enhancements

As GP is a stochastic process, convergence towards desired performance can take several generations. In order to ensure that convergence time does not strain energy and processing resources on smart phones, AGP includes two components, namely the helper and the supervisor, which enable the use of expert knowledge to apply constraints to AGP-generated programs.

4.3.1. Helper

The helper is called during the program generation process of any program builder. The developer can create one or several helpers for one application by implementing the AGP HelperInterface interface. Whenever a program

is generated, AGP will refer to the helper evaluate() method to check any correctness conditions that the program has to meet. For instance, in our geolocalization application (cf. Section 5.1.2), if a program doesn't call any location provider, we know that this program will be unable to locate the smartphone, so we can reasonably discard it without losing evaluation time, and therefore we save battery life.

4.3.2. Supervisor

The supervisor runs during the program interpretation. It can check constraints on-the-fly, and kill the Interpreter Shell if the program goes out of bounds. For instance, the supervisor can control execution time. If the program execution is too long, the supervisor will automatically kill the program.

4.4. Island Model

The Island Model is a parallelised genetic algorithm (GA) [37] which employs multiple, distinct subpopulations (islands) in order to promote the genetic diversity of the complete system. Individual genomes can migrate from one subpopulation to another in an attempt to mix diverse genetic material to promote genetic diversity and alleviate premature convergence.

The AGP framework creates a local population on each device, essentially an island with no migration inward (immigration) or outward (emmigration). However the inherent communications capabilities of mobile phones makes the Island Model an attractive extension to the AGP framework and also provides a mechanism for harnessing the collective resources of multiple devices for the benefit of all devices. These benefits can potentially manifest

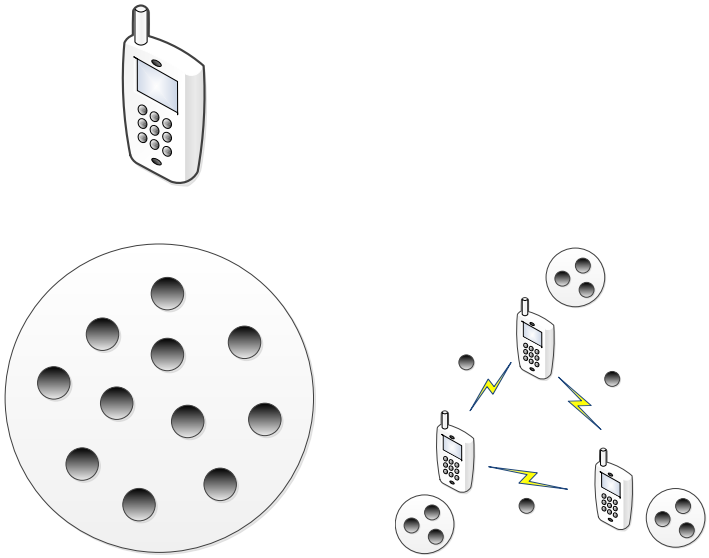

(a) Original AGP framework

(b) AGP framework with Island Model extension

Figure 3: Comparison of the original AGP framework and the Island Model extension. In the extension, each phone contains a local population that periodically exchanges genetic programs during migratory periods.

themselves as faster convergence times, more robust solutions and even more fit solutions. Figure 3 provides a synoptic comparison of the standalone AGP approach and the Island Model approach.

We extend AGP to support the island model by allowing the island running on each phone to send and receive migrant programs from and to its local population. Since smartphones incorporate short-range wireless communication interfaces, such as Bluetooth and Wifi, we use these interfaces for the transmission and reception of migrant programs. The full implementation details of AGP's features that support the island model are available online.

5. Performance Evaluation

This section demonstrates the capability of smart phone personalisation through online GP using 2 example applications.

5.1. Standalone Personalisation

We demonstrate standalone personalisation through AGP with applications on two different generations of Android devices: the early HTC Magic running Android 1.6 ; and the more capable Nexus S, embedding a dual-core processor with the recent Android Gingerbread (version 2.3).

5.1.1. Google Reader Personalisation

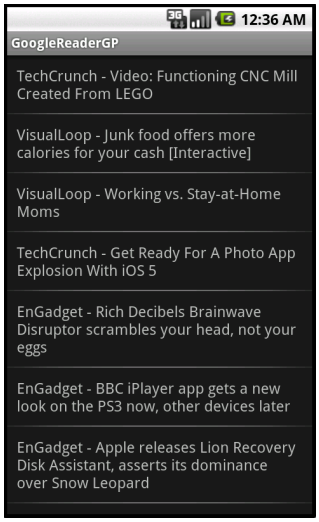
12:36 AM
GoogleReaderGP
TechCrunch - Video: Functioning CNC Mill Created From LEGO
VisualLoop - Junk food offers more calories for your cash [Interactive]
VisualLoop - Working vs. Stay-at-Home Moms
TechCrunch - Get Ready For A Photo App Explosion With iOS 5
EnGadget - Rich Decibels Brainwave Disruptor scrambles your head, not your eggs
EnGadget - BBC iPlayer app gets a new look on the PS3 now, other devices later
EnGadget - Apple releases Lion Recovery Disk Assistant, asserts its dominance over Snow Leopard

Figure 4: Screenshot of the GRGP interface

Google Reader is a web-based aggregator released and maintained by Google. It works as a RSS feed reader, allowing users to get the latest news from selected feeds. Many applications exist for extracting news feeds from

smartphones. Preferences for reading news on a smartphone may depend on the content type, particularly its readability on a mobile device with a small screen. For instance, it is easy to read short text news whereas it is uncomfortable to look for long articles, comics, infographics or flash animations. Moreover, such content might quickly deplete the user monthly capped data plan.

Available Google Reader apps on the applications market reports news from all the feeds to which the user has subscribed. Even though some feeds are not easily readable on a mobile device, the application doesn't learn from the user reading habits and keeps proposing items from all the feeds. Our Google Reader GP (GRGP) application takes advantage of the AGP framework to learn which feeds the user likes to read on their smartphone within a given context. The feeds are taken from the user Google Reader account.

The application is kept simple for demonstration purposes: whenever the user wants to get news, she asks for a news report which executes a GP program and returns the latest and unread news from feeds selected by the program.

Fitness definition. The fitness definition is based on two sub-fitness functions:

$$Fitness = Fitness_{count} \times Fitness_{clicked}$$

$$Fitness_{count} = \begin{cases} \frac{Disp.\ \ news}{Desired\ \ qty.} & Disp.\ news \leq Desired\ qty. \\ 1 & otherwise \end{cases}$$

$$Fitness_{clicked} = \left(\frac{Clicked\ \ news}{Displayed\ \ news} \right)$$

The news count refers to the minimum desired quantity of displayed news whenever the user requests a news report. This is a setting fixed by the user in the configuration menu. The clicked news corresponds to the quantity of news read over the amount of news given in the report. In other terms, it evaluates the interest of the user in the displayed news.

Population evolution. This experiment uses 5 program per population. Once every program in the population has been executed, they all will have received a fitness value. It is then time to generate a new population by using what has been learned. We use elitism by copying the fittest program into the new population. In GRGP, we use a typical roulette wheel selector. Thanks to selectors, evolution operators and tools embedded in AGP, we are able to clearly define our evolution strategy. Table 1 gives basic commands to call a basic strategy relying on mutations and cross-over done over the 3 best programs, which we refer to as highly ranked (HR) programs. We then perform mutations on the HR programs from the previous generation to generated 2 mutated programs, which are included in the new population. The final 2 programs in the new population are crossovers of HR programs in the previous population. The maximum number of epochs in the simulation was set to 100. We set the maximum program depth to 3 to limit the program bloat, which is common in GP. As we have chosen problems with limited complexity to demonstrate the capabilities of AGP, this control mechanism was sufficient to avoid any bloat in our experiments.

Table 1: Basic commands for defining evolution strategy

Action	**Command**
Set a wheel selector with the 3 best programs	setSelector(WHEEL,"HR", getnBest(3))
Set a wheel selector with the best program	setSelector(WHEEL,"Leader", getnBest(1))
Generate 2 mutated programs	generate("HR", MUTATION, 2)
Generate 2 programs from crossover	generate("HR", CROSSOVER, 2)
Copy the best program into the new population	generate("Leader", COPY, 1)
Get the new population	getOutputPopulation()

Results. We conduct our experiment with 7 sources from a real Google Reader subscription: 4 technology news websites (TechCrunch, TechLand, Engadget and Digital Trends), VisualLoop which gives fresh infographics, Break Videos for funny videos, and Business Green for latest green products. Even though the concerned user was interested in all those sources, he is used to read the technology news on his smartphone rather than the other sources providing longer articles or heavy media files, as they are not optimised for reading on smartphone. We set the target number of news stories to 10.

Figure 5 reports the average of displayed news per program over the first 30 generations, where we observe convergence. In order to provide a

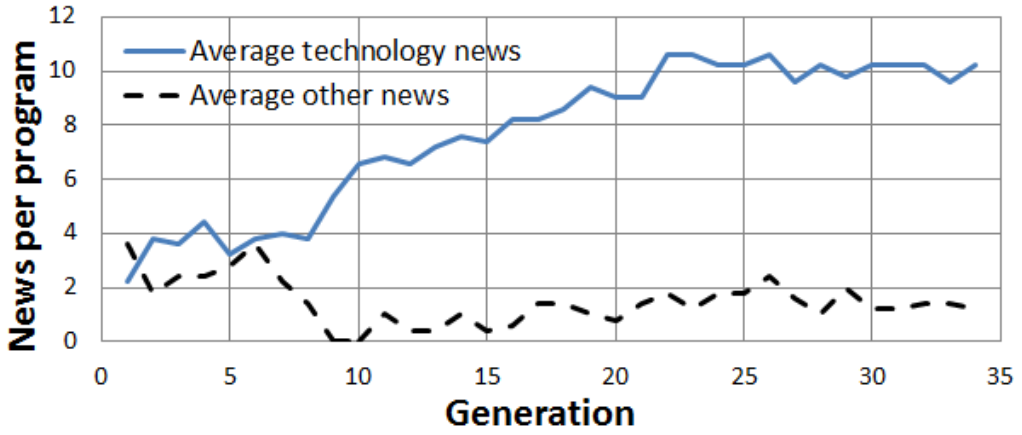

Figure 5: The evolution of the number of news stories displayed in the standalone Google Reader application.

more readable graph, we grouped the news feed in two sub-groups: the 4 technology news, and the others (VisualLoop, BreakVideos, and Business Green). After 6 generations, our Google Reader using AGP learned the user preference for the technology news. However, it doesn't eliminate completely the diversity and keeps proposing some news from the other feeds yet with a lower likelihood.

As expected, GRGP learns to provide the desired minimum quantity of news per report. Figure 5 confirms this convergence by looking at the sum of technology and other news.

The program fitness evolution indicates that our evolution strategy does its job, and leads to an increase of the elite program fitness. In Figure 6, the pool average represents the mean pool fitness (i.e. mean fitness of the 5 programs). Whenever the elite program gains in fitness, it subsequently leads to an increase of the pool average fitness.

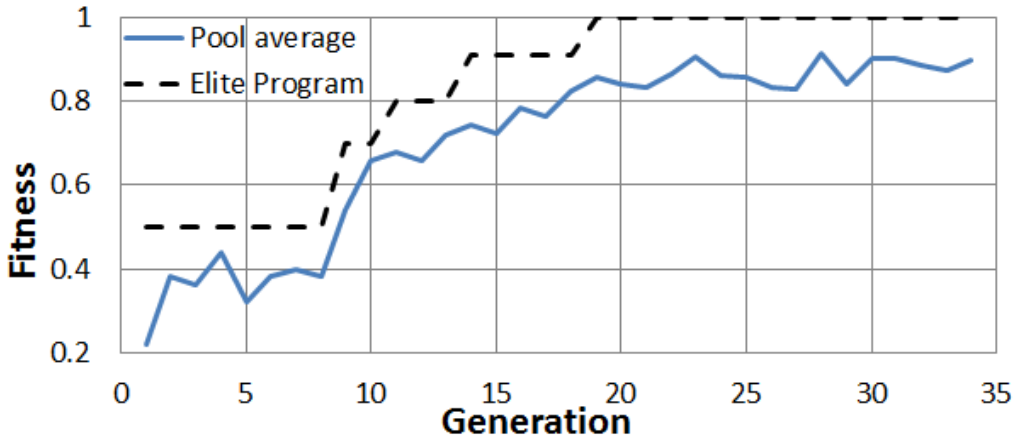

Figure 6: Program fitness evolution for the standalone Google Reader application.

5.1.2. Context-aware localization

We now turn our attention to adapting the localisation strategy running on the smart phone based on its context. Localization is a complex problem as there are multiple location services on a typical smart phone: GPS, which is energy-hungry yet accurate and only works outdoors; WiFi, which provides mid-level accuracy and energy consumption, and works when there are WiFi access points nearby; and cell-based localisation that relies on cellular phone towers, which is energy-efficient yet inaccurate. The various localization options on smartphones requires consideration of their availability and their energy/accuracy trade-off [23, 24].

As this problem depends on too many contextual constraints such as position, signal quality, and device energy profile, it is difficult to design a context free algorithm for selecting the best location service, particularly since the most accurate service depends on the user's current location and whether they are indoors or outdoors. We introduce a context-aware applica-

tion using AGP aimed to address this problem of selecting the best location service provider. The focus here is on AGP's ability to solve multi-objective problems online.

Fitness function definition. The fitness function used in our localization application reflects the common trade-off between energy and accuracy in the localization field [22]. We introduce two fitness metrics: the accuracy fitness and the energy fitness, which respectively quantify the accuracy and energy efficiency of the provided solution. As positions are dynamic, we evaluate fitness every second during the evaluation period (n seconds). By the end of the evaluation time, we use the average of these subfitnesses to give a fitness to the program. The overall fitness is obtained by multiplying these subfitnesses. We choose this option as it discards any solution which doesn't provide any accuracy or could deplete the battery, while maintaining simplicity for the demonstration purpose of this paper.

$$Fitness = \frac{\sum_{i=1}^{n} Fitness_{Acc}(i) \times Fitness_{En}(i)}{n} \tag{1}$$

Accuracy fitness. Most smart phones include several location providers, such as GPS, Cellular Network and Wi-Fi, that can be used alone or in combination. When the application is learning, the evaluation process keeps all the location providers on, which makes learning costly from an energy perspective. The application automatically picks the provider giving the most favorable accuracy. We call the output from this provider the best available position.

We use this best available position as a reference to attribute an accuracy

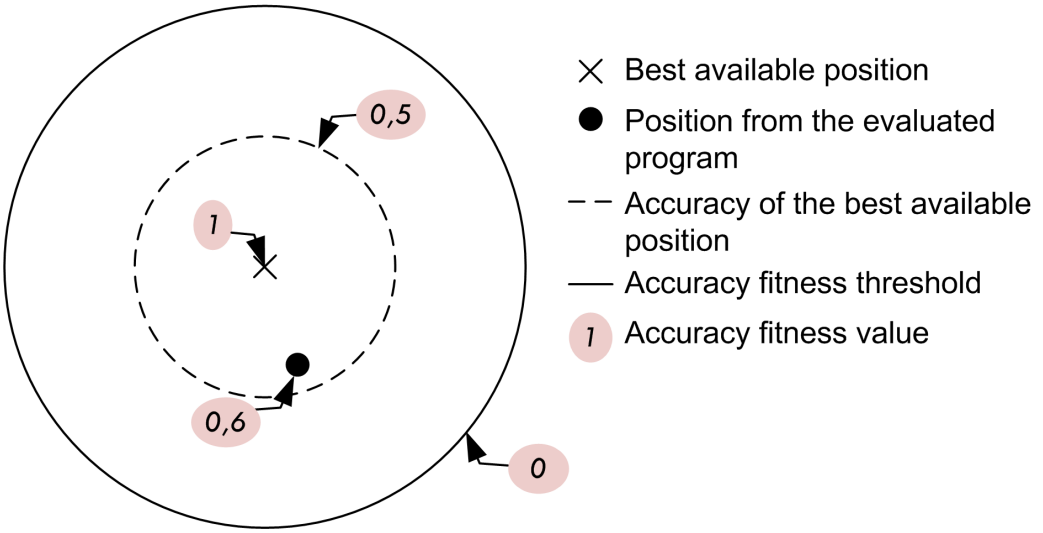

Figure 7: Fitness accuracy definition for context aware localisation

fitness to the position provided by the evaluated program (cf. Figure 7):
- if program position is within the accuracy of the best available position, we attribute an accuracy fitness following a linear rule from 1 to 0.5
- if program position is outside the former, but within a circle of twice the accuracy of the best available position (accuracy fitness threshold), we attribute an accuracy fitness following a linear rule from 0.5 to 0.

Energy fitness. We define energy fitness according to a basic rule: we want to provide the best possible localization accuracy without depleting the battery by the end of the day as we assume users can charge their phones at the end of each day. We consider an average 1400 mAh battery capacity for the paper. To achieve our goal, the average power consumption should not exceed 63 mA (= 1400 / 22 hours). We define a day as 22 hours because we consider the phone as plugged for 2 hours per day.

We use the Android PowerProfile class to estimate the power consumption per chip. As an internal OS function used to compute the battery usage, this class isn't accessible through the Android Developer API. However, it can be readily obtained through the Java reflection mechanism.

This enables AGP to assess the power cost of the evaluated program depending on the CPU usage and the chip used to locate the smartphone. For simplicity, we limit the energy fitness to a linear function, ranging from 1 for a idealistic case where the program doesn't cost any power to 0 for a program which requires more power than the one day energy budget.

Results. We conduct the experiment with populations of 12 programs, of which 1 is the fittest (elite) program from the previous generation, 4 are mutations, 5 are crossovers, and 2 are randomly generated programs. Mutations and crossover program are generated in the exact same way as for the Google Reader application, while the random programs are introduced to increase genetic diversity of the population. Because each program in this experiment requires much longer than the Google Reader Programs (for instance to acquire fixes from GPS), we set the experiment to terminate after 40 epochs. The program evaluation is limited to one minute to provide sufficient time for the GPS module to obtain a fix. Our function set contains general operators such as addition, multiplication, and other application-specific operators: functions to switch location providers such as GPS, Wi-Fi or Cellular Network. Figure 8 shows the framework ability to get a program localizing the smartphone. After a rough first solution, it converges to smarter programs able to provide more efficient and accurate solutions.

We also conduct an experiment to evaluate the benefits from the use of

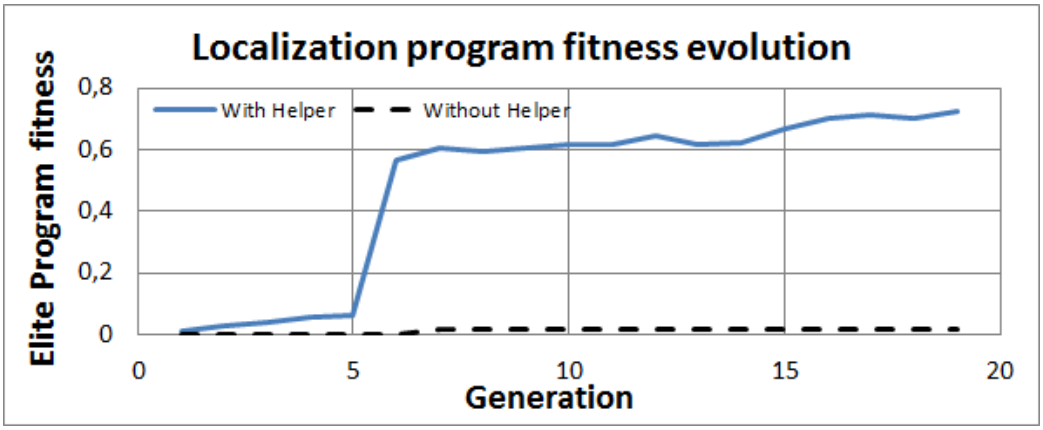

Figure 8: Fitness evolution for context-aware localisation

Helper. Populations generated with the Helper provide a working solution in the first generation, and quickly have satisfying programs. On the other hand, populations generated without the Helper are stuck with non-working programs (zero fitness) for several generations. Then, they only get a slow evolution. It is mainly due to many programs which make no sense: they don't switch on a location provider or don't call any latitude nor longitude update.

6. Collaborative Personalisation Experiments

Having demonstrated online smartphone personalisation through GP, we now evaluate collaborative personalisation on the same example applications.

6.1. Setup

Both example applications use two Samsung Google Nexus S smartphones running Android 2.3 Gingerbread. The population size for the collaborative personalisation experiments has 10 programs. The entire local subpopulation

is replaced each generation whereby the fittest program (elite) is retained between generations. The three HR programs are then mutated to create three new programs, and the remaining programs are generated using crossovers applied to the entire breeding pool. We use proportionate selection with respect to a program's fitness value, and we select emigrants at random without removing them from the local subpopulation. Immigrant programs are appended to the local subpopulation. Experiments are conducted using two different migration intervals (epoch length) i of 5 and 10 generations and three different migration rates r of 0.1, 0.2 and 0.3 of local population size — i.e. 10%, 20% and 30% of the local subpopulation are exchanged each migration period. We collect data about the genetic programs' fitness, size and tree depth for 20 generations and average it over 15 iterations to ensure accuracy and minimise the effect of the initial starting conditions.

The experiments use a migration schema, where genetic programs are randomly generated and added to the local subpopulation in lieu of the traditional migration between different subpopulations and smartphones. This enables the evaluation of the performance benefit versus energy consumed for the communication of genetic programs between phones. The genetic programs are randomly generated by an instance of the ProgramBuilder class and added to the subpopulation at a migration interval of $i = 5$ and a migration rate of $r = 0.3$.

Furthermore, for the Google Reader test application, we conduct further experiments using two different fitness landscape configurations: (1) a homogeneous fitness function over the two populations; and (2) a heterogeneous fitness function over the two populations. The differences between the fitness

functions for the Google Reader application are simply the news sources that the user was interested in — i.e. clicked. For example, in one population TechCrunch and Engadget news items are "clicked" and in the second population Break Videos and Digital Trends news items are "clicked" instead. Such differences in the fitness function are common in real world scenarios where each phone has a different inputs/constraints — e.g. different user preferences, different context or different phone type. Thus our results using the heterogeneous fitness function provide a more realistic evaluation of the Island Model's performance.

6.2. Results

Figure 9 shows the results from the experiments for the Google Reader application using a heterogeneous fitness function among the subpopulations and Figure 11 shows the results from the experiments conducted using the localisation example application. The figures show the experimental results for four of the six variations of the two migration parameters. In every tested configuration, the Island Model outperforms the non-migratory population model in terms of convergence time. For instance, in Figure 10(a) shows that the elite program with migration reaches a fitness of 90% by generation 4, while it takes until generation 12 to reach this fitness with the standalone model, representing a 66% improvement.

For the experiments using a homogeneous fitness function (c.f. Figure 10), we find that the fitness trajectory shows short exponential increases or spikes in the maximum fitness value in the generations immediately following a migration period, which can be seen at generation 4 for the experiments with i=5 and generation 9 for the experiments with i=10. This can be explained

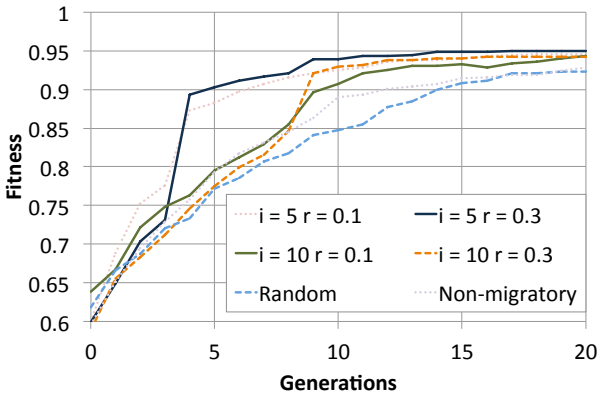

(a) Max Fitness

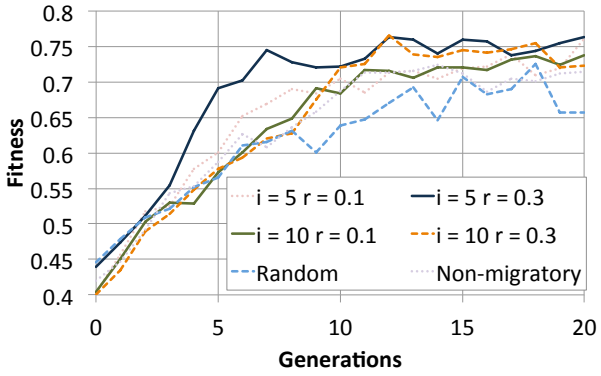

(b) Mean fitness

Figure 9: Fitness data for the Google Reader example application using a heterogeneous fitness function.

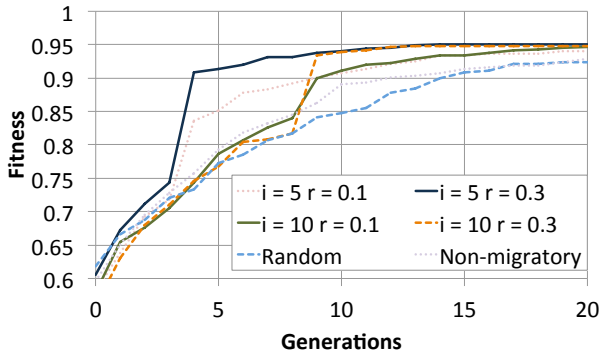

(a) Max Fitness

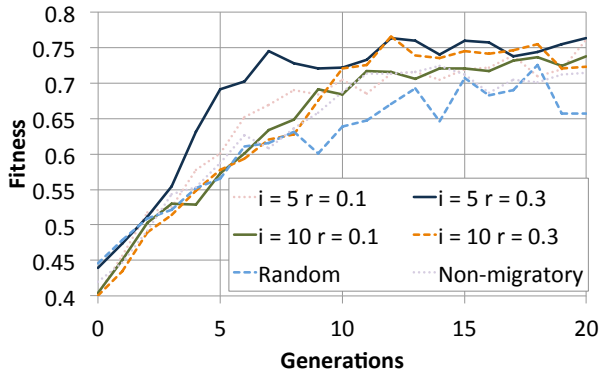

(b) Mean fitness

Figure 10: Fitness data for the Google Reader example application using a homogeneous fitness function.

by the building block hypothesis [34], where during the periods of isolated evolution the subpopulations evolve useful "building blocks". Subsequently when migration occurs between the two subpopulations, the good blocks in the migrant programs are assimilated within the local subpopulation's genetic pool leading to the construction of larger and better novel building blocks and resulting in fitter genetic programs.

We can see similar results for the Google Reader application using the Island Model and a heterogeneous fitness function, which display the same characteristic short exponential increases in the maximum fitness value after each migration period. However, the introduction of the migrants typically causes a temporary decrease in the average fitness followed by a rapid increase in the maximum fitness, supporting the results obtained by Martin et al [26], Mühlenbein [38] and Belding [36]. This can be attributed to the heterogeneous fitness function or fitness landscape, influencing each subpopulation to explore a different area within the genetic search space, helping to promote genetic diversity between subpopulations.

Whilst still somewhat noticeable in the Google Reader fitness data, the effect of the initial starting conditions on the results is more pronounced for the localisation application. This is highlighted by the different starting locations of the mean and maximum fitness values for the non-migratory model and each different Island Model configuration demonstrated by the higher standard deviation and standard error values for the fitness data. This effect could be exacerbated due smartphone's context relevant to the objective function not being completely static during the evolution and execution of the genetic programs — e.g. GPS reception and network reception and con-

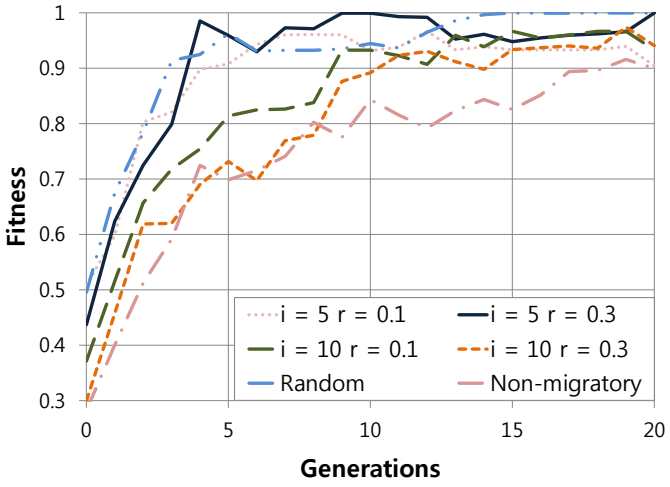

(a) Max fitness

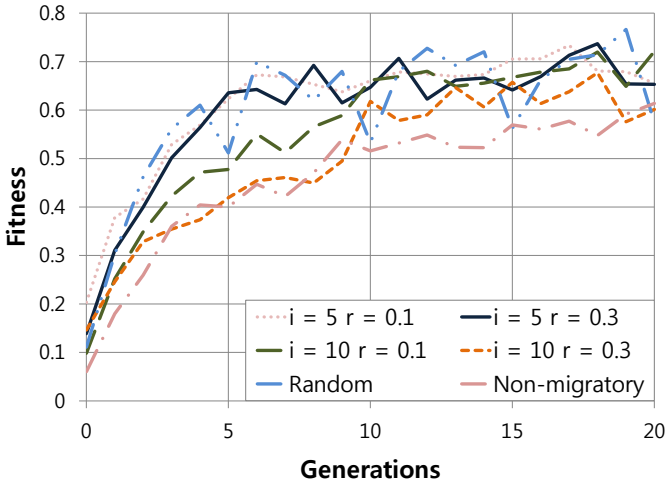

(b) Mean fitness

Figure 11: Fitness data for the localisation example application.

nection speed varying slightly over time and altering the fitness of executed genetic programs. It is recommended in future studies to use a larger sample size to mitigate this effect.

The results show there is no fitness benefit from the introduction of the randomly generated programs for the experiments using the Google Reader application. This contrasts with the results for the localisation application that demonstrate that the random programs have a potentially positive impact on the GA's performance. However, the average maximum fitness for the Google Reader application was much higher than for the non-migratory population model suggesting that the effect of the initial starting conditions is somewhat responsible for these results. We can also observe that there are significant dips in the average mean fitness immediately following the introduction of the randomly created programs. Followed by a slightly larger increase in the subpopulations mean fitness as the junk genetic material introduced during the migratory period is discarded through evolutionary processes. This can sometimes also result in a small increase to the subpopulations max fitness.

The inconsistent performance due to the introduction of random programs can be attributed to the effect of random sampling having a small, but finite, chance of introducing more fit genetic material than in the current population. However, the probability of this occurring is much less than the probability of randomly selected programs from other subpopulations having a good fitness and this probability will continue to diminish as the average subpopulation fitnesses increase. This causes some inconsistency in the performance improvement of the random programs over the standard

non-migratory population model. Nonetheless, these results are intriguing and accord with the intuition that any attempts to maintain or increase genetic diversity will typically have a positive influence on the population's evolution.

Based upon these results, an interesting idea for future work is the epigenetic selection of terminals, such as the strategy suggested in [42], to govern the creation of genetic programs. Essentially, grammars defined by probabilistic trees to select the function and terminal sets that are ultimately instantiated, could be dynamically evolved. For example, if some function A has terminals B and C as its children and is shown to have a high probability of a high fitness, then this bias could be used when generating new programs, such that there is a higher chance for any instance of the function A to have the terminals B and C as its children.

7. Discussion and Conclusion

The benefits of the Island Model over the standard non-migratory population model are well known [38, 36, 26, 39] and this is clearly demonstrated in the obtained results. There is a very obvious increase in learning speed that can be attributed to the effects of migration when employing the Island Model. This is observable in the evolutionary trajectories of the fitness scores in both the Google Reader and the localisation example applications. Notably, the mean (pool) fitness decreases immediately after a migratory period, due to the typically lower fitness of the immigrant programs, however on occasion, these dips are then followed by a short and sometimes exponential increase in the elite fitnesses of each subpopulation as newly diverse and

useful genetic material is incorporated into local subpopulation.

The experiments conducted with the Google Reader example application using a heterogeneous fitness function between subpopulations displayed a similar performance benefit to those using a homogeneous fitness function. Martin et al [26] suggest that this is due the different fitness functions promoting genetic diversity between the different subpopulations, consequently improving the exploitation and exploration of the genetic search space. This demonstrates that the heterogeneous nature of the pervasive computing environment — i.e. different users and preferences, different smartphone properties, helps rather than hinders the performance of the Island Model, allowing the Island Model to more efficiently leverage the available processing resources compared to the traditional non-migratory population model.

Finally, it must be noted that some privacy concerns have been raised about collaborative techniques, specifically about fine-grained localisation [35]. However, genetic programming is well placed to help preserve users' privacy by sharing logic containing variables and not needing to share the actual values of the variables (such as location, energy, etc). Furthermore, the logic shared may or may not represent the user's actual preference due to the stochastic selection of which programs are migrated and finally there is no need to reveal the source of the program, so ideally the programs would be broadcast anonymously.

Overall, the results of this study show the significant potential of the combination of the Island Model with the AGP framework for achieving collaborative learning of dynamic user personalisation in pervasive computing contexts. While we have tested the Island Model on only 2 phones and on

relatively simple applications as a proof of concept, we expect the benefits of parallel resources and improved population diversity as the number of participating islands (phones) increases to be able to address more complex applications and expedite convergence times [26, 41].

The symbolic nature of GP means that logic can be readily seeded and is an intuitive choice where control of device's resources is typically performed programatically. It should be noted, however, that adaptive behaviour may not be desirable for all interactions. For example, the user interface should have a consistent feel across applications and adhere to the device or OS recommended UI design recommendations. However, within this constraint, adaptive behaviour may provide a method to overcome consistently undesirable application (i.e. fixed logic) behaviour. The current configuration employs a fixed population structure which produces a number of likely low performing programs to be evaluated every generation even after the system has converged. This means that there will always be some residual unexpected or unpredictable behaviour, however this exploration aspect is a necessary evil in order to maintain adaptability to significant changes in user preferences and also to refine performance over time.

User satisfaction is mainly driven by AGP's capability to quickly adapt to context and provide suitable solutions. To achieve this goal, AGP needs a learning process which can have an efficient energy profile. As the end-user is also sensitive to early battery depletion, introducing a mechanism to tune this trade-off between battery and learning would help maintain satisfaction. An interesting direction for future work would be to implement a mechanism to configure AGP for intensive learning or a low energy consumption mode

depending on context.

Another area for future investigation is to extend the context aware localisation application and AGP to include other sensor streams that could aid localisation. For example, the developer could also introduce other providers such as a contact-logging beacon method [23, 24], or an accelerometer-assisted localisation algorithm [40]. Finally, an interesting direction for future work is to explore the use Pareto GP algorithms for multi-objective problems to evaluate its potential benefits over compounded fitness functions that combine several objectives.

In summary, this paper has proposed collaborative personalisation through online GP. We have demonstrated the capability for smart phone personalisation through online GP, and have quantified the benefits of collaborative personalisation with the Island Model. We believe this study represents a key first step towards versatile online personalisation in the growing smart phone market.

Acknowledgements

The authors would like to thank Brano Kusy for his valuable inputs in realising this work. This project was supported by the Sensors and Sensor Network Transformational Capability Platform at CSIRO.

References

[1] Lee, M., Cho S.: Interactive Differential Evolution for Image Enhancement Application in Smart Phone.In IEEE CEC (2012) 2411-2416

[2] Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R., Estrin, D.: Diversity in smartphone usage.In MobiSys(2010) 179-194

[3] Costa, P.D., Pires, L.F., Sinderen, M.V.: Designing a configurable services platform for mobile context-aware applications.Int. J. Pervasive Computing and Communications(2005) 13-24

[4] Bae, J., Lee, J.Y., Kim, B., Ryu, S.: Next Generation Mobile Service Environment and Evolution of Context Aware Services.In EUC(2006) 591-600

[5] Miele, A., Quintarelli, E., Tanca, L.: A methodology for preference-based personalisation of contextual data.In EDBT(2009) 287-298

[6] Sutton, Richard S. Temporal Credit Assignment in Reinforcement Learning (PhD thesis). (1984).

[7] J. J. HOPFIELD. Neural networks and physical systems with emergent collective computational abilities. Proc. NatL Acad. Sci. USA Vol. 79, pp. 2554-2558, April 1982

[8] Koza, J.R.: Genetic programming - on the programming of computers by means of natural selection.Complex adaptive systems(1993) 1-419

[9] Ventura, S., Romero, C., Zafra, A., Delgado, J.A., Hervs, C.: JCLEC: a Java framework for evolutionary computation.Soft Comput.(2008) 381-392

[10] McPhee, N. F., Hopper, N. J., Reierson, M. L.: Sutherland: An extensible object-oriented software framework for evolutionary computation.

In Genetic Programming 1998: Proceedings of the Third Annual Conference (1998) 241-253

[11] Gagn, C., Parizeau, M.: Open BEAGLE: A New Versatile C++ Framework for Evolutionary Computation.In GECCO Late Breaking Papers(2002) 161-168

[12] Ismail, I.A., Ramly, N.A.E., Abd-ElWahid, M.A., ElKafrawy, P.M., Nasef, M.M.: Genetic Programming Framework for Fingerprint Matching.CoRR(2009) 316-321

[13] Torres, R.D.S., Falco, A.X., Gonalves, M.A., Papa, J.P., Zhang, B., Fan, W., Fox, E.A.: A genetic programming framework for content-based image retrieval.Pattern Recognition(2009) 283-292

[14] Valencia, P., Lindsay, P., Jurdak, R.: Distributed genetic evolution in WSN.In IPSN(2010) 13-23

[15] Lenaerts, T., Manderick, B.: Building a Genetic Programming Framework: The Added-Value of Design Patterns.In EuroGP(1998) 196-208

[16] Lacerda, A., Cristo, M., Gonalves, M.A., Fan, W., Ziviani, N., Ribeiro-Neto, B.A.: Learning to advertise.In SIGIR(2006) 549-556

[17] Weise, T., Geihs, K.: DGPF: An Adaptable Framework for Distributed Multi-Objective Search Algorithms Applied to the Genetic Programming of Sensor Networks. In BIOMA'06, October 910, 2006, pages 157-166, Ljubljana, Slovenia. (2006)

[18] Korpipaa, P., Malm, E-J., Rantakokko, T., Kyllonen, V., Kela, J., Mantyjarvi, J., Hakkila, J., Kansala, I.: Customizing User Interaction in Smart Phones. In IEEE Pervasive Computing 5, 3 (July 2006), 82-90.

[19] onX project, `https://www.onx.ms/`

[20] Launch Center Pro, `http://appcubby.com/launch-center/`

[21] Cotillon, A., Valencia, P., Jurdak, R.: Android Genetic Programming Framework.In EuroGP(2012) 13-24

[22] Lin, K., Kansal, A., Lymberopoulos, D., Zhao, F.: Energy-accuracy trade-off for continuous mobile device location.In MobiSys(2010) 285-298

[23] Jurdak, R., Corke, P., Dharman, D., Salagnac, G.: Adaptive GPS duty cycling and radio ranging for energy-efficient localization.In SenSys(2010) 57-70

[24] R. Jurdak, P. Corke, A. Cotillon, D. Dharman, C. Crossman, and G. Salagnac, "Energy-efficient Localisation: GPS Duty Cycling with Radio Ranging," ACM Transactions on Sensor Networks, Vol. 9, Iss. 2, May 2013.

[25] Folino, G., Spezzano, G.: P-CAGE: An Environment for Evolutionary Computation in Peer-to-Peer Systems.In EuroGP(2006) 341-350

[26] W. N. Martin, J. Lienig and J. P. Cohoon: Island (Migration) Models: Evolutionary Algorithms Based on Punctuated Equilibria, in Handbook of Evolutionary Computation, Oxford University Press, (1997)

[27] Herdy, Michael, Evolution strategies with subjective selection, Parallel Problem Solving from Nature Lecture Notes in Computer Science, Eds. Voigt, Hans-Michael and Ebeling, Werner and Rechenberg, Ingo and Schwefel, Hans-Paul, Springer, 1141 (1996), 22-31

[28] Kosorukoff, A., Human based genetic algorithm, IEEE International Conference on Systems, Man & Cybernetics, (2001), 3464

[29] K. Sims, Artificial Evolution for Computer Graphics, Computer Graphics (Proc. Siggraph), Vol. 25, No. 4, July (1991), 319-328

[30] Craig Caldwell, Victor S. Johnston: Tracking a Criminal Suspect Through "Face-Space" with a Genetic Algorithm. ICGA (1991): 416-421

[31] Yao, Xin. "Evolutionary artificial neural networks." International journal of neural systems 4.03 (1993): 203-222.

[32] Faulkner, M., Olson, M., Chandy, R., Krause, J., Chandy, K., Krause, A.: The next big one: Detecting earthquakes and other rare events from community-based sensors. In: Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on. pp. 13–24 (april 2011)

[33] Weiss, G.M., Lockhart, J.W.: The impact of personalization on smartphone-based activity recognition. In: Proceedings of the Activity Context Representation Workshop. pp. 98–104 (2012)

[34] Holland, J.: Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence. Bradford (1975)

[35] Wicker, S.B.: The loss of location privacy in the cellular age. Commun. ACM 55(8), 60–68 (Aug 2012)

[36] Belding, T.C.: The distributed genetic algorithm revisited (1995)

[37] Cantu-Paz, E.: A summary of research on parallel genetic algorithms (1995)

[38] Mühlenbein, H.: Evolution in time and space - the parallel genetic algorithm. In: Foundations of Genetic Algorithms. pp. 316–337. Morgan Kaufmann Publishers Inc. (1991)

[39] Turner, J.: Genetic algorithms with punctuated equilibria: Parallelization and analysis (1998)

[40] R. Jurdak, B. Kusy, and A. Cotillon, "Group-based Motion Detection for Energy-efficient Localization," Journal of Sensor and Actuator Networks. 1(3):183-216, October 2012.

[41] Tatsuya Niwa and Masaru Tanaka. Analysis on the Island Model Parallel Genetic Algorithms for the Genetic Drifts. In Selected papers from the Second Asia-Pacific Conference on Simulated Evolution and Learning on Simulated Evolution and Learning (SEAL'98), Bob McKay, Xin Yao, Charles S. Newton, Jong-Hwan Kim, and Takeshi Furuhashi (Eds.). Springer-Verlag, London, UK, UK, 349-356. (1998)

[42] R. P. Salustowicz and J. Schmidhuber. Probabilistic incremental program evolution. Evolutionary Computation, 5(2):123141. (1997)